

WHAT IS CLAIMED IS:

1. A method for job management in an HPC environment comprising:

5 determining an unallocated subset from a plurality of HPC nodes, each of the unallocated HPC nodes comprising an integrated fabric;

selecting an HPC job from a job queue; and

executing the selected job using at least a portion of the unallocated subset of nodes.

10

2. The method of Claim 1, wherein selecting the HPC job comprises selecting the HPC job from the job queue based on priority, the selected job comprising dimensions not greater than a topology of the unallocated
15 subset.

3. The method of Claim 2, wherein selecting the HPC job from the job queue based on priority comprises:

sorting the job queue based on job priority;

20 selecting a first HPC job from the sorted job queue;

determining dimensions of the first HPC job with the topology of the unallocated subset; and

in response to the dimensions of the first HPC job being greater than the topology of the unallocated
25 subset, selecting a second HPC job from the sorted job queue.

4. The method of Claim 2, wherein the dimensions of the first HPC job are based, at least in part, on one
30 or more job parameters and an associated policy.

5. The method of Claim 2, further comprising:
dynamically allocating a job spare from the
unallocated subset based, at least in part, on the
dimensions of the HPC job; and

5 wherein executing the selected job comprises
executing the selected job using the dynamically
allocated job spare.

6. The method of Claim 1, the plurality of HPC
10 nodes comprising a first plurality and the method further
comprising:

determining that dimensions of the selected job are
greater than a topology of the first plurality;

15 selecting one or more HPC nodes from a second
plurality, each of the second HPC nodes comprising an
integrated fabric; and

adding the selected second HPC nodes to the
unallocated subset to satisfy the dimensions of the
selected job.

20

7. The method of Claim 6, further comprising
returning the second HPC nodes to the second plurality.

8. The method of Claim 1, further comprising;
25 determining that a second HPC job that was executing
on a second subset in the plurality of HPC nodes has
failed;

adding the second subset to the unallocated subset;
and

30 adding the failed job to the job queue.

9. Software for job management in an HPC environment operable to:

determine an unallocated subset from a plurality of HPC nodes, each of the unallocated HPC nodes comprising
5 an integrated fabric;

select an HPC job from a job queue; and
execute the selected job using at least a portion of the unallocated subset of nodes.

10 10. The software of Claim 9, wherein the software operable to select the HPC job comprises software operable to select the HPC job from the job queue based on priority, the selected job comprising dimensions not greater than a topology of the unallocated subset.

15 11. The software of Claim 10, wherein the software operable to select the HPC job from the job queue based on priority comprises software operable to:

sort the job queue based on job priority;
20 select a first HPC job from the sorted job queue;
determine dimensions of the first HPC job with the topology of the unallocated subset; and

in response to the dimensions of the first HPC job being greater than the topology of the unallocated
25 subset, select a second HPC job from the sorted job queue.

12. The software of Claim 10, wherein the dimensions of the first HPC job are based, at least in
30 part, on one or more job parameters and an associated policy.

13. The software of Claim 10, further operable to:
dynamically allocate a job spare from the
unallocated subset based, at least in part, on the
dimensions of the HPC job; and

5 wherein the software operable to execute the
selected job comprises software operable to execute the
selected job using the dynamically allocated job spare.

14. The software of Claim 9, the plurality of HPC
10 nodes comprising a first plurality and the software
further operable to:

determine that dimensions of the selected job are
greater than a topology of the first plurality;

15 select one or more HPC nodes from a second
plurality, each of the second HPC nodes comprising an
integrated fabric; and

add the selected second HPC nodes to the unallocated
subset to satisfy the dimensions of the selected job.

20 15. The software of Claim 14, further comprising
returning the second HPC nodes to the second plurality.

16. The software of Claim 9, further operable to:
determine that a second HPC job that was executing
25 on a second subset in the plurality of HPC nodes has
failed;

add the second subset to the unallocated subset; and
add the failed job to the job queue.

17. A system for job management in an HPC environment comprising:

a plurality of HPC nodes, each node including an integrated fabric; and

5 a management node operable to:

determine an unallocated subset from the plurality of HPC nodes;

select an HPC job from a job queue; and

10 execute the selected job using at least a portion of the unallocated subset of nodes.

18. The system of Claim 17, wherein the management node operable to select the HPC job comprises the management node operable to select the HPC job from the
15 job queue based on priority, the selected job comprising dimensions not greater than a topology of the unallocated subset.

19. The system of Claim 18, wherein the management
20 node operable to select the HPC job from the job queue based on priority comprises the management node operable to:

sort the job queue based on job priority;

select a first HPC job from the sorted job queue;

25 determine dimensions of the first HPC job with the topology of the unallocated subset; and

in response to the dimensions of the first HPC job being greater than the topology of the unallocated subset, select a second HPC job from the sorted job
30 queue.

20. The system of Claim 18, wherein the dimensions of the first HPC job are based, at least in part, on one or more job parameters and an associated policy.

5 21. The system of Claim 18, further operable to:
dynamically allocate a job spare from the unallocated subset based, at least in part, on the dimensions of the HPC job; and

10 wherein the management node operable to execute the selected job comprises the management node operable to execute the selected job using the dynamically allocated job spare.

15 22. The system of Claim 17, the plurality of HPC nodes comprising a first plurality and the management node further operable to:

determine that dimensions of the selected job are greater than a topology of the first plurality;

20 select one or more HPC nodes from a second plurality, each of the second HPC nodes comprising an integrated fabric; and

add the selected second HPC nodes to the unallocated subset to satisfy the dimensions of the selected job.

25 23. The system of Claim 22, the management node further operable to return the second HPC nodes to the second plurality.

24. The system of Claim 17, the management node further operable to:

determine that a second HPC job that was executing on a second subset in the plurality of HPC nodes has
5 failed;

add the second subset to the unallocated subset; and
add the failed job to the job queue.